

attribute of R , and Q an SQL query that contains $R.A$ in the SELECT clause. $SIT(R.A|Q)$ is the statistic for attribute A on the result of the executing query expression Q . Q is called the generating query expression of $SIT(R.A|Q)$. This definition can be extended for multi-attribute statistics. Furthermore, the definition can be used as the basis for extending the CREATE STATISTICS statement in SQL where instead of specifying the table name of the query, more general query expression such as a table valued expression can be used.

In U.S. Patent Application Serial No. 10/191,822, incorporated herein by reference in its entirety, the concept of SITs was introduced. A particular method of adapting a prior art query optimizer to access and utilize a preexisting set of SITs for cost estimation was described in detail in this application, which method is summarized here briefly as background information.

Referring to Figure 2, the query optimizer examines an input query and generates a query execution plan that most efficiently returns the results sought by the query in terms of cost. The cost estimation module and its imbedded cardinality estimation module can be modified to utilize statistics on query expressions, or intermediate tables, to improve the accuracy of cardinality estimates.

In general, the use of SITs is enabled by implementing a wrapper (shown in phantom in Figure 2) on top of the original cardinality estimation module of the RDBMS. During the optimization of a single query, the wrapper will be called many times, once for each different query sub-plan enumerated by the optimizer. Each time the query optimizer invokes the modified cardinality estimation module with a query plan, this input plan is transformed by the wrapper into another one that exploits SITs. The

is now US Patent # 6,947,927
file
1/19/2006

Best Available Copy

cardinality estimation module uses the input plan to arrive at a potentially more accurate cardinality estimation that is returned to the query optimizer. The transformed query plan is thus a temporary structure used by the modified cardinality and is not used for query execution.

is not US patent # 6,947,927

According to the embodiment described in application serial number 10/191,822, the transformed plan that is passed to the cardinality estimation module exploits applicable SITs to enable a potentially more accurate cardinality estimate. The original cardinality estimation module requires little or no modification to accept the transformed plan as input. The transformation of plans is performed efficiently, which is important because the transformation will be used for several sub-plans for a single query optimization.

1/19/2006

In general, there will be no SIT that matches a given plan exactly. Instead, several SITs might be used for to some (perhaps overlapping) portions of the input plan. The embodiment described in application serial number 10/191,822 integrates SITs with cardinality estimation routines by transforming the input plan into an equivalent one that exploits SITs as much as possible. The transformation step is based on a greedy procedure that selects which SITs to apply at each iteration, so that the number of independence assumptions during the estimation for the transformed query plan is minimized. Identifying whether or not a SIT is applicable to a given plan leverages materialized view matching techniques as can be seen in the following example.

is not US patent # 6,947,927

In the query shown in Figure 3(a) $R \bowtie S$ and $R \bowtie T$ are (skewed) foreign-key joins. Only a few tuples in S and T verify predicates $\sigma_{S.a < 10}(S)$ and $\sigma_{T.b > 20}(T)$ and most tuples in R join precisely with these tuples in S and T . In the absence of SITs,

Available Copy

Therefore error values must be estimated using efficient and coarse mechanisms. Existing information such as system catalogs or characteristics of the input query can be used but not additional information created specifically for such purpose.

IS NOS US Patent # 6,947,927

1/19/2006

Application serial number 10/191,822 introduced an error function, $nInd$, that is simple and intuitive, and uses the fact that the independence assumption is the main source of errors during selectivity estimation. The overall error of a decomposition is defined as $S = Sel_{R_1}(P_1|Q_1) \cdot \dots \cdot Sel_{R_n}(P_n|Q_n)$ when approximated, respectively, using $\mathcal{H}_{R_1}(A_1|Q'_1), \dots, \mathcal{H}_{R_n}(A_n|Q'_n)$ ($Q'_i \subseteq Q_i$), as the total number of predicate independence assumptions during the approximation, normalized by the maximum number of independence assumptions in the decomposition (to get a value between 0 and 1). In symbols, this error function is as follows:

$$nInd(\{Sel_{R_i}(P_i|Q_i), \mathcal{H}_{R_i}(A_i|Q'_i)\}) = \frac{\sum_i |P_i| \cdot |Q_i - Q'_i|}{\sum_i |P_i| \cdot |Q_i|}$$

Each term in the numerator represents the fact that P_i and $Q_i - Q'_i$ are independent with respect to Q_i , and therefore the number of predicate independent assumptions is $|P_i| \cdot |Q_i - Q'_i|$. In turn, each term in the denominator represents the maximum number of independence assumptions when $Q'_i = \emptyset$, i.e. $|P_i| \cdot |Q_i|$. As a very simple example, consider $S = Sel_R(R.a < 10, R.b > 50)$ and decomposition $S = Sel_R(R.a < 10 | R.b > 50) \cdot Sel_R(R.b > 50)$. If base table histograms $H(R.a)$ and $H(R.b)$ are used, the error using $nInd$ is $\frac{1 \cdot (1 - 0) + 1 \cdot (0 - 0)}{1 \cdot 1 + 1 \cdot 0} = 1/1 = 1$, i.e., one out of one independence assumptions (between

Best Available Copy